arXiv:1512.00540v1 [cs.DC] 2 Dec 2015

Dynamic Multiple-Message Broadcast: Bounding Throughput in the Affectance Model

Dariusz R. Kowalski[1], Miguel A. Mosteiro[2], and Kevin Zaki[2]

[1] Univ. of Liverpool, Dept. of Computer Science, Liverpool, UK.
D.Kowalski@liv.ac.uk
[2] Kean University, Dept. of Computer Science, Union, NJ.
{mmosteir,zakike}@kean.edu

Abstract. We study a dynamic version of the Multiple-Message Broadcast problem, where packets are continuously injected in network nodes for dissemination throughout the network. Our performance metric is the ratio of the throughput of such protocol against the optimal one, for any sufficiently long period of time since startup. We present and analyze a dynamic Multiple-Message Broadcast protocol that works under an affectance model, which parameterizes the interference that other nodes introduce in the communication between a given pair of nodes. As an algorithmic tool, we develop an efficient algorithm to schedule a broadcast along a BFS tree under the affectance model. To provide a rigorous and accurate analysis, we define two novel network characteristics based on the network topology and the affectance function. The combination of these characteristics influence the performance of broadcasting with affectance (modulo a logarithmic function). We also carry out simulations of our protocol under affectance. To the best of our knowledge, this is the first dynamic Multiple-Message Broadcast protocol that provides throughput guarantees for continuous injection of messages and works under the affectance model.

1 Introduction

We study the dynamic Multiple-Message Broadcast problem in wireless networks under the *affectance* model. This model subsumes many communication-interference models studied in the literature, such as Radio Network (cf., [7]) and models based on the Signal to Interference and Noise Ratio (SINR) (cf. [20, 32]). The notion of affectance was first introduced in [20] in the context of link scheduling in the more restricted SINR model of wireless networks, in an attempt to formalize the combination of interferences from a subset of links to a selected link under the SINR model. Later on, other realizations of affectance were defined and abstracted as an independent model of interference in wireless networks [24, 25]. The conceptual idea of this model is to parameterize the interference that transmitting nodes introduce in the communication between a given pair of nodes.

Our results. In the *dynamic Multiple-Message Broadcast* problem considered in this work, packets arrive at nodes in an online fashion and need to be delivered to all nodes in the network. We are interested in the *throughput*, i.e., the number of packets delivered in a given period of time. In particular, we measure *competitive throughput* of *deterministic distributed algorithms* for the dynamic Multiple-Message Broadcast problem. We analyse our algorithms in the (general) affectance model, in which there is a given undirected communication graph G of n nodes and diameter D, together with the affectance function $a(\cdot)$ of nodes of distance at least 2 on each of the communication links. The affectance function has a *degradation parameter* α, being a distance after which the affectance is negligible. Our contribution is two fold.

First, we introduce new model characteristics — based on the underlying communication network and the affectance function — called *maximum average tree-layer affectance* (denoted by K) and *maximum*

A preliminary version of this work has appeared in [28]. The differences with respect to that version are detailed in Section B in the Appendix.

path affectance (denoted by M), see Section 2 for the definitions, and show how they influence the time complexity of broadcast. More precisely, if one uses a BFS tree that minimizes the product $M \cdot (K + M/\log n)$[3] of the two above characteristics, then a single broadcast can be done in time $D + O(M(K + M/\log n)\log^2 n)$, cf., Corollary 1 in Section 3.

Second, we extend this method of analysis to a dynamic packet arrival model and the Multiple-Message Broadcast problem, and design a new algorithm reaching competitive throughput of $\Omega(1/(\alpha K \log n))$. In particular, in the Radio Network model it implies a competitive throughput of $\Omega(1/(\log^2 n))$. For details, see Section 4. Our deterministic results are existential, that is, we show the existence of a deterministic schedule by applying a probabilistic argument to a protocol that includes a randomized subroutine for layer to layer dissemination. Given that we measure competitive throughput in the limit, preprocessing (communication infrastructure setup, topology information dissemination, etc.) can be carried out initially without asymptotic impact. Thus, the protocol presented is distributed, and it works for *every* network after learning its topology. The protocol can also be applied to mobile networks, if the movement is slow enough to recompute the structure. Our rigorous asymptotic analysis is further complemented by simulations under the affectance model, c.f., Section 5.

To the best of our knowledge, ours is the first work on the dynamic Multiple-Message Broadcast problem in wireless networks under the general affectance model.

Previous and related work. There is a rich history of research on broadcasting dynamically arriving packets on a *single-hop radio network*, also called a *multiple access channel*. Most of the research focused on *stochastic arrivals*, cf., a survey by Chlebus [9]. In the remainder of this paragraph, we focus on the online adversarial packet arrival setting. Bender et al. [5] studied stability, understood as throughput being not smaller than the packet arrival rate, of randomized backoff protocols on multiple access channels in the *queue-free model*, in which every packet is handled independently as if it has been a standalone station (thus avoiding queuing problems). Kowalski [27] considered a dynamic broadcast on the channel in the setting where packets could be combined in a single message, which again avoids various important issues related with queuing. Anantharamu et al. [3] studied packet latency of deterministic dynamic broadcast protocols for arrival rates *smaller* than 1. Stability, understood as bounded queues, of dynamic deterministic broadcast on multiple access channels against adversaries bounded by arrival rate 1 was studied by Chlebus et al. [11], and for arrival rates smaller than 1 by Chlebus et al. [12]. In particular, in [11] a protocol Move-big-to-front (MBTF) was designed, achieving stability but not fairness (as both these properties are impossible to achieve simultaneously); we use this algorithm as a subroutine in our dynamic Multiple-Message Broadcast protocol. A follow-up work [6] delivered a distributed online algorithm SCAT and showed that it could be only by a linear factor worse, in terms of the buffer size, than any offline solution against any arrival pattern.

In *multi-hop* Radio Networks, the previous research concentrated on time complexity of single instances (i.e., from a single source) of broadcast and multi-message broadcast. For directed networks, the best deterministic solution is a combination of the $O(n \log n \log\log n)$-time algorithm by De Marco [16] and the $O(n \log^2 D)$-time algorithm by Czumaj and Rytter [14]. In undirected networks, the best up to date deterministic broadcast in $O(n \log(n/D))$ rounds was given by Kowalski [27]. The lower bounds for deterministic broadcast in directed and undirected radio networks are $\Omega(n \log(n/D))$ [13] and $\Omega(n \log_D n)$ [29], respectively. Deterministic multi-message broadcast, group communication and gossip were also considered (again, in a single instance). Chlebus et al. [10] showed a $O(k \log^3 n + n \log^4 n)$ time deterministic multi-broadcast algorithm for k packets in undirected radio networks. Single broadcast can be done optimally in $\Theta(D \log(n/D) + \log^2 n)$, as proved in [2, 31] (lower bounds) and in [14, 29] (matching upper bound). Bar-Yehuda et al. [4], and recently Khabbazian and Kowalski [26] and Ghaffari et al. [19], studied ran-

[3] Throughout, we denote $\log_2$ simply as $\log$, unless otherwise stated.

domized multi-broadcast protocols; the best results obtained for k-sources single-instance multi-broadcast is the amortized $O(\log \Delta)$ rounds per packet w.h.p. in [26], where Δ is the maximum node degree, and $O(D + k \log n + \log^2 n)$ w.h.p. to broadcast the k packets, for settings with known topology in [19]. For the same problem, Ghaffari et al. showed a throughput upper bound of $O(1/\log n)$ for any algorithm in [18]. Although this bound is worst-case, it can be compared with our $1/O(\alpha K \log n)$ that applies even under affectance.

Chlebus et al. [11] gave various deterministic and randomized algorithms for group communication, all of them being only a small polylogarithm away of the corresponding lower bounds on time complexity.

In the *SINR model*, single-hop instances of broadcast in the ad-hoc setting were studied by Jurdzinski et al. [22, 23] and Daum et al. [15], who gave several deterministic and randomized algorithms working in time proportional to the diameter multiplied by a polylogarithmic factor of some model parameters. In the SINR model with restricted sensitivity, so called weak-sensitivity device model, Jurdzinski and Kowalski [21] designed an algorithm spanning an efficient backbone sub-network, that might be used for efficient implementation of multi-broadcast.

The *generalized affectance* model was introduced and used only in the context of one-hop communication, more specifically, to link scheduling by Kesselheim [24]. He also showed how to use it for dynamic link scheduling in batches. This model was inspired by the affectance parameter introduced in the more restricted SINR setting [20]. They give a characteristic of a set of links, based on affectance, that influence the time of successful scheduling these links under the SINR model. In our paper, we generalize this characteristic, called the maximum average tree-layer affectance, to be applicable to multi-hop communication tasks such as broadcast, together with another characteristic, called the maximum path affectance. For details see Section 2.

2 Preliminaries

Model. We study a model of network consisting of n ***nodes***, where communication is carried out through radio ***transmissions*** in a shared channel. Time is discretized in a sequence of time slots $1, 2, \ldots$, which we call the ***global time***. The network is modeled by the underlying ***connectivity graph*** $G = \{V, E\}$, where V is the set of nodes and E the set of ***links*** among nodes. A link $\ell \in E$ between two nodes $u, v \in V$ is the ordered pair (u, v) modeling that a transmission from u may be received by v. The network is assumed to be connected but ***multihop***. That is, any pair of nodes may communicate, possibly through multiple hops.

Messages to be broadcast to the network through radio transmissions are called ***packets***. Packets are ***injected*** at nodes at the beginning of time slots, and each time slot is long enough to transmit a packet to a neighboring node. Any given node can either transmit or listen (in order to receive, if possible) in a time slot, but not both.

Interference on a link due to transmissions from other nodes is modeled as ***affectance***. We use a model of affectance that subsumes other interference models [4], such as the Radio Network model [7] and the SINR model [20]. Specifically, we realize affectance as a matrix A of size $|V| \times |E|$ where $A(u, (v, w))$ quantifies the interference that a transmitting node $u \in V$ introduces to the communication through link $(v, w) \in E$. We do not restrict ourselves to any particular affectance function, as long as its effect is additive. That is, denoting $a_{V'}((v, w))$ as the affectance of a set of nodes $V' \subseteq V$ on a link $(v, w) \in E$, and $a_{V'}(E')$ as the

[4] In preliminary work [28], we studied a different model of affectance. The details are included in Section B in the Appendix.

affectance of a set of nodes $V' \subseteq V$ on a set of links $E' \subseteq E$, it is

$$a_{V'}((v,w)) = \sum_{\substack{u \in V' \\ u \neq v}} A(u,(v,w))$$
$$a_{V'}(E') = \sum_{(v,w) \in E'} a_{V'}((v,w)).$$

Under the affectance model, we define a ***successful transmission*** as follows. For any pair of nodes $u, v \in V$ such that $(u,v) \in E$, a transmission from u is received at v in a time slot t if and only if: u transmits and v listens in time slot t, and $a_{\mathcal{T}(t)}((u,v)) < 1$, where $\mathcal{T}(t) \subseteq V$ is the set of nodes transmitting in time slot t (notice that the definition of a does not include the affectance of u on (u,v)). The event of a non-successful transmission, that is when the affectance is at least 1, is called a ***collision***. We assume that a node listening to the channel cannot distinguish between a collision and background noise present in the channel in absence of transmissions.

The affectance model defined subsumes other interference models. For instance, for the Radio Network model, the affectance matrix is

$$A(w,(u,v)) = \begin{cases} 0 \text{ if } w = u \text{ or } ((w,v) \notin E \text{ and } w \neq v), \\ 1 \text{ otherwise.} \end{cases}$$

On the other hand, for the SINR model in [20], the affectance matrix is

$$A(w,(u,v)) = \begin{cases} 0 & \text{if } w = u, \\ \frac{P/d_{uv}^{\alpha}}{P/(\beta' d_{uv}^{\alpha}) - N} & \text{otherwise.} \end{cases}$$

Where P is the transmission power level, N is the background noise, β' denotes an upper bound on the signal to interference-plus-noise ratio such that a message cannot be successfully received, d_{uv} is the euclidean distance between nodes u and v, and α denotes the path-loss exponent. (Refer to Section A in the Appendix for a proof.)

Communication task. Under the above model, we study the ***Multiple-Message Broadcast*** problem defined as follows. Starting at time slot 1, packets are dynamically injected by an exogenous entity into some of the network nodes, called ***source nodes***. The computing task is to disseminate those injected packets throughout the network. The set of all source nodes is denoted as $S \subseteq V$. After a packet has been received by all the nodes in the network, we say that the packet was ***delivered***. The injections are adversarial, that is, packets can be injected at any time slot at any source node, but the injections are limited to be feasible. We say that an injection is ***feasible*** if there exists an optimal algorithm OPT such that the ***latency*** (i.e., the time elapsed from injection to delivery) of each packet is bounded for OPT. Given that at most one packet may be received by a node in each time slot, and that all nodes must receive the packet to be delivered, feasibility limits the adversarial injection rate to at most 1 packet per time slot injected in the network. The goal is to find a ***broadcasting schedule***, that is, a temporal sequence of transmit/not-transmit states for each node, so that packets are delivered. We denote the period of time since a packet is transmitted from the source until its delivery as the ***length*** of the schedule.

Performance metric. We evaluate the ***ratio*** of the performance of a distributed online algorithm ALG against an optimal algorithm OPT. For one hop networks it is known [11] that no protocol is both ***stable*** (i.e.,

bounded number of packets in the system at any time) and ***fair*** (i.e., every packet is eventually delivered). For multihop networks the same result holds as a natural extension of the single hop model. Thus, instead of further limiting the adversary (beyond feasibility) to achieve either stability or bounded latency, our goal is to prove a lower bound on the ***competitive throughput***, for any sufficiently long prefix of time slots since global time 1. Specifically, we want to prove that there exists a function f, possibly depending on network parameters, such that

$$\lim_{t\to\infty} \frac{d_{ALG}(t)}{d_{OPT}(t)} \in \Omega(f),$$

where $d_X(t)$ is the number of packets delivered to all nodes by algorithm X until time slot t.

Network characterization. We characterize a network by its ***affectance degradation distance***, which is the number of hops α such that the affectance of nodes of distance at least α to a given link is "negligible", that is, zero. Additionally, we characterize the network with two measures of affectance based on broadcast trees [5], as follows. Given a network with a set of nodes V including a source node s, consider a Breadth First Search (BFS) tree T rooted at s. For any $d = 0, 1, 2, \ldots$, let $V_d(T)$ be the set of all nodes at (shortest) distance d from s. Based on this tree, we define the ***maximum average tree-layer affectance*** as

$$K(T,s) = \max_d \max_{V' \subseteq V_d(T)} \frac{1}{|L(V')|} a_{V'}(L(V')) \ ,$$

where $L(V')$ is the set of tree links between V' and nodes at distance $d+1$ of the source. Intuitively, $K(T,s)$ indicates what might be the worst affectance to overcome when trying to broadcast from one layer of T to another. We also define the ***maximum path affectance*** as

$$M(T,s) = \max_{p \in P(T)} \sum_{(u,v)\in p} a_{V_{d(u)}(T)}((u,v)) \ ,$$

where $d(u)$ is the distance from node u to s, and $P(T)$ is the set of paths root-to-leaf in T (i.e., a set of sets of links), where a ***path root-to-leaf*** is the standard notion of a set of links $\{(s,x_1),(x_1,x_2),(x_2,x_3),\ldots,(x_{k-1},x_k)\}$ such that x_k is a leaf. Intuitively, $M(T,s)$ indicates what is the worst affectance when trying to pipeline packets through a path down the tree. In the rest of the paper, the specific tree and source node will be omitted when clear from context.

3 A Broadcast Tree

In this section, we show a broadcasting schedule that, under the affectance model, disseminates a packet held at a source node to all other nodes. The schedule is defined constructively with a protocol that uses randomization, thus providing only stochastic guarantees. Given that the protocol is Las Vegas, the construction also proves the existence of a deterministic broadcasting schedule.

First, we detail the construction of a ranked tree spanning the network rooted at the source node that will be used to define the broadcasting schedule that we detail afterwards. The construction borrows the idea in [17] of defining some nodes as *fast* and others as *slow* based on rank. However, our rank is a consequence of affectance rather than Radio Network collisions, and it is defined to schedule transmissions downwards the tree only, rather than both directions. Moreover, our definition of the fast node sets, the slot reservation, and the contention resolution protocol are also different. The following notation will be used.

[5] The second characterization was presented differently in the conference version of this work. The details are included in Section B in the Appendix.

Given a tree $T(s) \subseteq E$ rooted at $s \in V$, spanning a set of network nodes V with set of links E, let $d(v)$ be the ***distance*** in hops from a node $v \in V$ to the root of $T(s)$, let $p(\ell)$ and $c(\ell)$ be the parent and child nodes of link $\ell \in T(s)$ respectively, and let $D(T(s))$ be the maximum distance in $T(s)$ from any node to the root s. Additionally, a ***rank*** (a number in $\mathbb{N}$) will be assigned to each node. Let $r(u)$ be the rank of node $u \in V$, let $R(T(s))$ be the maximum rank in the tree, and let $F_d^r(T(s)) = \{u | u \in V_d \wedge r(u) = r \wedge \exists v \in V_{d+1} : ((u,v) \in T(s) \wedge r(v) = r)\}$, that is, the set of nodes of rank r at distance d from the root that have a child with the same rank. Let a node $v \in V$ be called ***fast*** if it belongs to the set $F_{d(v)}^{r(v)}(T(s))$, and ***slow*** otherwise. The sets F_d^r are called ***fast node sets*** whereas the set containing all slow nodes is called ***slow node set***. In the above notation, the specific tree parameter and/or source node will be omitted when clear from the context.

Given a graph G and a source node $s \in S$, consider the following construction of a ***Low-Affectance Broadcast Spanning Tree (LABST)***. Let T_{min} be the BFS tree that minimizes the following polynomial on the affectance measures. Letting $\mathcal{T}$ be the class of all BFS trees with source s, it is

$$\forall T \in \mathcal{T} : M(T_{\min}, s) \left(\frac{M(T_{\min}, s)}{\log n} + K(T_{\min}, s) \right) \leq M(T, s) \left(\frac{M(T, s)}{\log n} + K(T, s) \right).$$

Then, using Algorithm 1, we define a rank on each node, that is, transform $T_{\min}$ into a LABST T, to avoid links between nodes of the same rank with big affectance.

In brief, the transformation is the following (refer to Algorithm 1). Initially, the rank of all nodes is set to 1, and the fast node sets are initialized in Lines 2 to 8. Then, for each distance d upwards the tree, two phases are executed as follows.

In a first phase (Lines 10 to 15), the rank of all nodes at distance d is updated if necessary. That is, for each increasing rank r, and for each link ℓ such that the parent node u is located at distance d (hence, child node at distance $d+1$) and parent and child nodes have rank r, check the affectance on link ℓ from other rank-r-distance-d nodes with a rank-r child. If this affectance is at least 1, increase the rank of u, and remove u from the fast node set F_d^r since its rank is not r anymore. Notice that u had the maximum rank among its children because it was in a fast node set, but now has a rank bigger than any of its children. Hence, u is now in the slow node set.

In the second phase (Lines 16 to 22), the rank of all the ancestors (distance $< d$) is updated so that the rank of a node at distance d equals the maximum rank (not necessarily unique) among its children (that is, ranks are monotonically non-decreasing upwards). While computing the rank, all fast node sets F_d^r are also updated.

The broadcasting schedule is defined using the LABST T obtained. Being a radio-broadcast network, transmissions might be received using other links or time slots, but the LABST and broadcasting schedule defined provide the communication guarantees. Each node follows certain broadcasting schedule, but using only time slots reserved for itself. Then, for each node $v \in V$, if v is fast, it uses each time slot t such that $t \equiv d(v) + 2h(R(T) - r(v)) \pmod{2hR(T)}$, where $h = \max\{3, \alpha\}$ and α is the affectance degradation distance. The purpose of lower bound h to 3 is to isolate affectance among neighboring layers, as in the Radio Network model. Otherwise, if v is slow, it uses each time slot t such that $t \equiv d(v) + h \pmod{2h}$. Notice that this schedule separates transmissions that occur in the same time step as follows. For any pair of slow nodes, or pair of fast nodes, they are either at the same distance from the source or they are separated by at least $2h$ hops from each other. For any pair of one slow and one fast node, they are separated by at least h hops. The reason to lower bound h by 3 is to avoid unnecessary interference between links separated by one hop.

The ***broadcasting schedule*** for fast nodes is simple: upon receiving a packet for dissemination, transmit in the next time slot reserved. For slow nodes, the schedule is determined by a randomized contention

Algorithm 1: LABST construction.

```
1  T ← T_min
2  foreach distance d = 0, 1, 2, ..., D(T) do
3      foreach rank r = 1, 2, ..., D(T) + 1 do                    // R(T) ≤ D(T) + 1
4          F_d^r ← ∅
5  foreach u ∈ V do
6      r(u) ← 1                                                   // initially, all nodes have rank 1
7      if u is not a leaf in T then
8          F_{d(u)}^1 ← F_{d(u)}^1 ∪ {u}
9  foreach distance d = D(T) − 1, ..., 2, 1, 0 do
10     foreach rank r = 1, 2, ..., D(T) + 1 do
11         foreach node u such that u ∈ F_d^r do
12             foreach link ℓ such that p(ℓ) = u do
13                 if u ∈ F_d^r and a_{F_d^r}(ℓ) ≥ 1 then
14                     r(u) ← r(c(ℓ)) + 1
15                     F_d^r ← F_d^r \ {u}                        // u is now slow
16     foreach distance d' = d − 1, ..., 2, 1, 0 do
17         foreach node u ∈ V_{d'} do                             // update ranks and sets
18             F_{d'}^{r(u)} ← F_{d'}^{r(u)} \ {u}
19             r_max ← max_{w∈V_{d'+1}, (u,w)∈T} r(w)              // max rank of children of u
20             if r(u) ≤ r_max then
21                 r(u) ← r_max
22                 F_{d'}^{r(u)} ← F_{d'}^{r(u)} ∪ {u}
```

resolution protocol that can be run in the reserved time slots. The protocol is simple: upon receiving a packet for dissemination, each slow node transmits repeatedly with probability $1/(4K(T_{\min}, s))$, until the packet is delivered.

In the rest of this section, we bound the length of the broadcasting schedule. The following upper bound will be used.

Lemma 1. *The maximum rank of a LABST T with source node s is*

$$R(T) \leq \lceil M(T_{\min}, s) \rceil.$$

Proof. Consider the construction of a LABST T in Algorithm 1. Consider any path from root to leaf in $T_{\min}$. Because initially all nodes have rank 1 (cf. Line 6), by definition of $M(T, s)$, the total affectance on this path is at most $M(T_{\min}, s)$. Each time that a node in such path increases its rank in Line 14, the node becomes slow in Line 15, and will not be fast again because the rank updates are carried level-by-level upwards the tree (cf. Line 9). Thus, after the transformation, the claimed bound holds because a value ≥ 1 is reduced from the total affectance due to fast nodes in the path (cf. Line 13).

Theorem 1. *For any given network of n nodes with a source node, diameter D, and affectance degradation distance α, there exists a broadcasting schedule of length at most*

$$D + 2h\lceil M(T_{\min})\rceil(\lceil M(T_{\min})\rceil + 16K(T_{\min}) \ln n),$$

where $h = \max\{3, \alpha\}$.

Proof. First we show that the broadcasting schedule is correct. Consider any pair of nodes $u, v \in V$ transmitting in the same time slot. If $d(u) = d(v)$ and they are both fast nodes with the same rank, the affectance on each other's links is low by definition of the LABST. If $d(u) = d(v)$ and they are both slow nodes, the contention resolution protocol will disseminate the packet to the next layer. Otherwise, given the slot reservation, it is either $|d(u) - d(v)| = 2h$ if u and v are both fast or both slow, or $|d(u) - d(v)| = h$ if one is slow and the other fast. Given that $h \geq \alpha$, the affectance on each other's links is negligible.

To prove the schedule length, consider any path p from root to leaf in the LABST T. The path p can be partitioned into consecutive maximal subpaths according to rank. In each maximal subpath $p' \in p$ of consecutive nodes of the same rank, the first node may have to wait up to $2hR(T)$ slots for the next reserved time slot, but after that all nodes except the last one transmit in consecutive time slots. Given that there are at most $R(T)$ such maximal subpaths and that their aggregated length is at most $D(T)$, the schedule length in the fast nodes of path p is at most $D(T) + 2hR(T)^2 \leq D + 2hR(T)^2$, where the latter inequality holds because T is a BFS tree.

Consider now any link $\ell \in p$ where the rank changes, that is $r(p(\ell)) \neq r(c(\ell))$ and $p(\ell) \in S_{d(p(\ell))} \subseteq V_{d(p(\ell))}$. Recall that the schedule in such link is defined by a randomized contention resolution protocol where each node transmits with probability $1/(4K(T_{\min}))$, where

$$K(T_{\min}) = \max_{d} \max_{V' \subseteq V_d(T_{\min})} \frac{1}{|L(V')|} a_{V'}(L(V')),$$

where $L(V')$ is the set of BFS tree links between V' and nodes at distance $d+1$ of the source, and $V_d(T_{\min})$ is the set of nodes at distance d from the source in $T_{\min}$. For a probability of transmission

$$q \leq \frac{1}{4 \max_{S \subseteq V_{d(p(\ell))}} a_S(L(S))/|L(S)|},$$

it was proved in [25] that the probability that there is still some link in S where no transmission was successful after $4c \ln |V_{d(p(\ell))}|/q$ time slots running Algorithm 1 in [25], is at most $|V_{d(p(\ell))}|^{1-c}$, $c > 1$. Given that $1/(4K(T_{\min}))$ verifies such condition, we know that after

$$16cK(T_{\min}) \ln |V_{d(p(\ell))}| \leq 16cK(T_{\min}) \ln n$$

(reserved) time slots, the transmission in link ℓ has been successful with positive probability. Given that there are at most $R(T) - 1$ links where the rank changes, using the union bound, we know that after $(R(T) - 1)16cK(T_{\min}) \ln n$ (reserved) time slots all slow nodes have delivered their packets with some positive probability, which shows the existence of a deterministic schedule of such length[6]. The time slots reserved for slow nodes appear with a frequency of $2h$. Thus, the schedule length in the slow nodes of path p is at most $2h(R(T) - 1)16cK(T_{\min}) \ln n \leq 32hR(T)K(T_{\min}) \ln n$, for $c = R(T)/(R(T) - 1)$.

Adding both schedule lengths we have

$$D + 2hR(T)^2 + 32hR(T)K(T_{\min}) \ln n$$

Replacing the bound on $R(T)$ in Lemma 1, the claim follows.

For networks with affectance degradation distance $\lceil \log n \rceil$, Theorem 1 yields the following corollary.

Corollary 1. *For any given network of $n \geq 8$ nodes, diameter D, and affectance degradation distance $\lceil \log n \rceil$, there exists a broadcasting schedule of length*

$$D + O\left(\log^2 n M(T_{\min}) \left(\frac{M(T_{\min})}{\log n} + K(T_{\min})\right)\right).$$

For comparison, for less contentious networks where affectance at more than one hop is not present (Radio Network model), using a GBST a broadcast schedule of length $D + O(\log^3 n)$ was shown in [17] and of length $O(D + \log^2 n)$ was proved in [30].

4 A Dynamic Multiple-Message Broadcast Protocol

In this section, we present our Multiple-Message Broadcast protocol and we bound its competitive throughput. The protocol uses the LABST[7] presented in Section 3.[8] The intuition of the protocol is the following. Each source node has a (possibly empty) queue of packets that have been injected for dissemination. Then, starting with an arbitrary source node $s \in S$ with "large enough" number of packets in its queue, packets are disseminated through a LABST rooted at s. If the number of packets in the queue of s becomes "small", s stops sending packets and, after some delay to clear the network, another source node $s' \in S$ starts disseminating packets through a LABST rooted at s'. The procedure is repeated following the order of a list of source nodes, which is dynamically updated according to queue sizes to guarantee good throughput. Packets from any given source are pipelined with some delay to avoid collisions and affectance. Being a radio broadcast network, packets might be received earlier than expected using links or time slots other than those defined by the LABST. If that is the case, to guarantee the pipelining, nodes ignore those packets.

[6] In settings with collision detection and where the affectance on any given link is $O(n)$, a big enough constant $c > 1$ yields a randomized protocol that succeeds with probability $1 - 1/n$.

[7] We refer to the tree and the broadcast schedule indistinctively.

[8] Any broadcast schedule that works under the affectance model could be used.

The following notation will be also used. The LABST rooted at $s \in S$ is denoted as $T(s)$. We denote the length of the broadcast schedule (time to deliver to all nodes) from s as $\Delta(s)$, and $\Delta = \max_{s \in S} \Delta(s)$. Let the pipeline delay (the time separation needed between consecutive packets to avoid collisions and affectance) from s be $\delta(s)$, and $\delta = \max_{s \in S} \delta(s)$. Given a node $i \in S$ and time slot t, the length of the queue of i is denoted $\ell(i,t)$. Let the length of all queues at time t be $\ell(t) = \sum_{i \in S} \ell(i,t)$. We say that, at time t, a node i is ***empty*** if $\ell(i,t) < \Delta$, ***small*** if $\Delta \leq \ell(i,t) < n\Delta$, and ***big*** if $\ell(i,t) \geq n\Delta$.

Consider the following ***Multiple-Message Broadcast Protocol***.

1. For each source node $s \in S$ define a LABST rooted at s.
2. Define a Move-big-to-front (MBTF) list [11] of source nodes, initially in any order. According to this list, source nodes circulate a token. While being disseminated, the token has a time-to-live counter of Δ, maintained by all nodes relaying the token. A source node s receiving the token has to wait for the token counter to reach zero before starting a new transmission. Let the time slot when the counter reaches zero be t. Then, node s does the following depending on the length of its queue.
   (a) If s is empty at t, it passes the token to the next node in the list. We call this event a ***silent round***.
   (b) If s is small at t, it broadcasts Δ packets pipelining them in intervals of δ slots. After δ more slots, it passes the token to the next node in the list.
   (c) If s is big at t, it moves itself to the front of the list. We call this event a ***discovery***. Then, s broadcasts packets pipelining them in intervals of δ slots as long as it is big, but a minimum of Δ packets. With the first of these packets s broadcasts the changes in the list. δ more slots after transmitting these packets, it passes the token to the next node in the list.

The following theorem shows an upper bound on the number of packets in the system at any time, which allows to prove the competitive throughput of our protocol. The proof structure is similar to the proof in [11] for MBTF, but many details have been redone to adapt it to a multihop network.

Theorem 2. *For any given network of n nodes, at any given time slot t of the execution of the Multiple-Message Broadcast protocol defined, the overall number of packets in queues is $\ell(t) < (t\delta/(1+\delta)) + 2\Delta n^2$.*

Proof. For the sake of contradiction, assume that there exists a time t such that the overall number of packets in the system is $\ell(t) \geq (t\delta/(1+\delta)) + 2\Delta n^2$. The number of packets in queues at the end of any given period of time is at most the number of packets in queues at the beginning of such period, plus the number of time slots when no packet is delivered, given that at most one packet is injected in each time slot. We arrive to a contradiction by upper bounding the number of time slots when no packet is delivered within a conveniently defined period before t. Consider the period of time T such that

$$\ell(t-T) \leq n^2\Delta + \frac{(t-T)\delta}{1+\delta} \tag{1}$$

$$\forall t' \in [t-T, t] : \ell(t') \geq n^2\Delta \tag{2}$$

$$\ell(t) \geq (t\delta/(1+\delta)) + 2\Delta n^2 \tag{3}$$

From now on, the analysis refers to the period of time T. We omit to specify it for clarity. Let $C \subseteq S$ be the set of nodes that are big at some point. Due to the pigeonhole principle and Equation (2), we know that for each time slot there is at least one big source node. In other words, the token cannot be passed throughout the whole list without at least one discovery. As a worst case, assume that only nodes in C have packets to transmit. For each node $i \in C$, the token has to be passed through at most $|S \setminus C| \leq n - |C|$ nodes that are not in C before i is discovered, because after i is discovered no node in $S \setminus C$ will be before i in the list.

Hence, there are at most $|C|(n-|C|)$ silent rounds, each of length Δ for token pass. So, due to passing the token through nodes in $S \setminus C$, there are at most $|C|(n-|C|)\Delta$ time slots when no packet is delivered.

We bound now the time slots when no packet is delivered due to passing the token through nodes in C before being discovered for the first time. Consider any given node $i \in C$. The argument is similar to the previous case. Any other node $j \in C$ that is discovered before i is moved to the front of the list. If i is going to be before j in the list later, it is not going to happen before i is discovered for the first time. Then, before i is discovered, it may hold the token at most $|C|-1$ times. As a worst case, assume that for each of these times i is empty. Hence, there are at most $|C|(|C|-1)$ silent rounds, each of length Δ for token pass. So, due to passing the token through nodes in C before being discovered, there are at most $|C|(|C|-1)\Delta$ time slots when no packet is delivered.

It remains to bound the time slots when no packet is delivered due to pipelining and passing the token through nodes in C after being discovered. Consider any given node $i \in C$ after being discovered. If i is big during the rest of T, it broadcasts packets pipelining them in intervals of δ slots. If instead i becomes small during T, i will have Δ packets to transmit for at least $n-1$ times that holds the token afterwards before becoming empty, because right after becoming small it has at least $(n-1)\Delta$ packets in queue. And there are at most $n-1$ nodes in C that will not be behind i in the list until i becomes big again. Hence, i always has Δ packets to transmit after being discovered the first time. After becoming small, i has to pass the token to the next node in the list introducing a delay of Δ. As a worst case scenario, we assume that upon each discovery of each node $i \in C$, only Δ packets are broadcast before passing the token. Then, for each Δ packets delivered, there are at most $\Delta + \Delta(\delta-1) = \Delta\delta$ time slots when no packet is delivered, over a period of $\Delta + \Delta\delta = \Delta(1+\delta)$ time slots. Because C is the set of nodes that are discovered in T, we can bound the number of batches of Δ packets delivered in T by $\lfloor T/(\Delta(1+\delta))\rfloor \leq T/(\Delta(1+\delta))$. Then, there are at most $T\Delta\delta/(\Delta(1+\delta)) = T\delta/(1+\delta)$ time slots when no packet is delivered due to nodes in C after being discovered.

Combining these bounds with Equation (1), we have that there are at most

$$\begin{aligned} n^2\Delta + \frac{(t-T)\delta}{1+\delta} + |C|(n-|C|)\Delta + |C|(|C|-1)\Delta + \frac{T\delta}{1+\delta} \\ = n^2\Delta + \frac{t\delta}{1+\delta} + \Delta|C|(n-1) \\ < \frac{t\delta}{1+\delta} + 2\Delta n^2 \end{aligned}$$

time slots when no packet is delivered. Which is a contradiction.

Lemma 2. *There exists a Multiple-Message Broadcast protocol that achieves a competitive throughput of at least*

$$\lim_{t\to\infty} \frac{1}{1+\delta} - \frac{2\Delta n^2}{t}.$$

Proof. A packet is delivered when it has been received by *all* nodes. The optimal algorithm delivers at most one packet per time slot, since any given node can receive at most one packet per time slot. Additionally, the injection is limited to be feasible, that is, there must exist an optimal algorithm OPT such that the latency of each packet is bounded for OPT. Thus, at most one packet may be injected in each time slot. Then, the competitive throughput is at least

$$\lim_{t\to\infty} \frac{d_{ALG}(t)}{d_{OPT}(t)} \geq \lim_{t\to\infty} \frac{t - nd_{ALG}(t)}{t},$$

where $nd_{ALG}(t)$ is the max number of packets that could not be delivered by ALG by time t. Using the bound in Theorem 2 we have that

$$\lim_{t\to\infty} \frac{d_{ALG}(t)}{d_{OPT}(t)} \geq \lim_{t\to\infty} \frac{t - (t\delta/(1+\delta)) - 2\Delta n^2}{t}$$
$$\geq \lim_{t\to\infty} \frac{1}{1+\delta} - \frac{2\Delta n^2}{t}.$$

The following theorem shows our main result.

Theorem 3. *For any given network of n nodes, diameter D, and affectance degradation distance α, there exists a Multiple-Message Broadcast protocol that achieves a competitive throughput of at least*

$$\lim_{t\to\infty} \frac{1}{1+\delta} - \frac{2\Delta n^2}{t}.$$

Where $\Delta \leq D + 2h\lceil M\rceil(\lceil M\rceil + 16K\ln n)$, $h = \max\{3,\alpha\}$, $K = \max_{s\in S} K(T_{\min}(s), s)$, $M = \max_{s\in S} M(T_{\min}(s), s)$, and $\delta \leq 16hK\ln n$.

Proof. The length $\Delta(s) \leq \Delta$ of the broadcast schedule in a LABST rooted at s is given in Theorem 1. With respect to $\delta(s) \leq \delta$, as explained in the proof of Theorem 1, slow nodes at distance d from the root deliver a packet to the next node in a path of a LABST $T(s)$ within $16cK(T_{\min}(s))\ln|V_d|$ with positive probability for any $c > 1$. This shows the existence of a deterministic schedule of that length. Additionally, packets must be separated by at least $\max\{3,\alpha\}$ to avoid collisions and affectance from nodes at different distances from the source (see the proof of Theorem 1 for further details). Then, it is $\delta(s) = \max\{3,\alpha\}16K(T_{\min}(s))\ln n$, for $c = \ln n/\ln|V_d|$. Replacing, the claim follows.

The above theorem yields the following corollary that provides intuition.

Corollary 2. *For any given network of n nodes, diameter D, and affectance degradation distance α, there exists a Multiple-Message Broadcast protocol such that the competitive throughput converges to*

$$\frac{1}{O(\alpha K\log n)},$$

where $K = \max_{s\in S} K(T_{\min}(s), s)$.

To evaluate these results, it is important to notice that the competitive throughput bound was computed against a theoretical optimal protocol that delivers one packet per time slot, which is not possible in practice in a multi-hop network. For comparison, instantiating our interference model in the Radio Network model (no affectance), using the WEB protocol [8] for slow transmissions our Multiple-Message Broadcast protocol can be shown to converge to $1/O(\log^2 n)$. Furthermore, for single-instance multi-broadcast in Radio Network, Ghaffari et al. showed in [18] a throughput upper bound of $O(1/\log n)$ for any algorithm. Although this bound is worst-case, it can be compared with our $1/O(\alpha K\log n)$ that applies even under affectance.

5 Simulations

We carried out simulations of our Multiple-Message Broadcast protocol under the affectance model.

For each of three values of $n = 16, 32, 64$, we produced three types of input networks to study the impact of high, medium and low interference. Namely, (1) a complete bipartite graph of n nodes split evenly (high interference), (2) two overlapped trees obtained running BFS on a connected random graph from different nodes (medium interference), and (3) a single path of nodes (low interference). For these inputs, we set the affectance degradation distance α to $\max\{1, \sqrt{\log n}\}$, $\max\{1, (\log n)/2\}$, and $\max\{1, \log n\}$ respectively. Recall that our model of general affectance comprises any interference effect (as long as it is additive). Hence, rather than restricting to a specific model for our simulations (such as SINR or Radio Network models which are geometric), we produced a less restrictive affectance matrix that comprises only the effect of the distance in hops from the interfering node to the receiver of the link. Specifically, for each input graph $G = \{V, E\}$, we computed an affectance matrix as follows. Let $d(i, k)$ be the shortest distance in hops from node i to node k. Then, for each node $i \in V$ and directed link $(j, k) \in E$,

$$A(i,(j,k)) = \begin{cases} 0 & \text{if } d(i,k) \geq \alpha \\ 1 & \text{if } d(i,k) = 0 \\ 1/d(i,k)^2 & \text{if } 0 < d(i,k) < \alpha \end{cases}$$

To simulate our Multiple-Message Broadcast protocol, we first computed the broadcast schedule from each node (although later we select only a subset of nodes for injection). For each node i, we computed a BFS tree rooted at i, we ranked the tree according to affectance, and we computed the average tree-layer affectance. To ensure polynomial running time, we computed *some* BFS tree rooted at each node i, rather than aiming for the tree that minimizes the polynomial on the affectance metrics (cf. Section 3) which would require to compute *all* BFS trees rooted at i. For the same reason, we computed $K(T, i) = \max_d a_{V_d(T)} L(V_d(T))/|L(V_d(T))|$, rather than aiming in each layer for the subset of tree links that maximizes the average affectance, which would require to compute the average affectance of *all* subsets of links.

For the input networks produced uniformly at random the above simplification should not have a significant impact in performance. Also, taking the tree-layer average affectance of all nodes in the layer guarantees that the contention resolution protocol used for slow nodes disseminates the packet to the next layer also under the affectance model, within the time bounds specified in the proof of Theorem 1, since interference from *all* nodes in the layer is taken into account.

With respect to the set of source nodes S, each node was chosen to be a source at random with probability $1/3$. For each source s, we computed the length of the schedule (time to deliver to all other nodes) as $\Delta(s) = D(T) + 2hR(T)^2 + 32hR(T)K(T)\ln n$, and the pipeline delay (time separation needed between consecutive packets) as $\delta(s) = 16hK(T)\ln n$. Notice that we used $D(T)$ and $R(T)$ rather than their bounds in Theorem 1, since for the simulations we know their values. Finally we computed $\Delta = \max_{s\in S}\{\Delta(s)\}$ and $\delta = \max_{s\in S}\{\delta(s)\}$.

The queue of one source node was initialized to $2\Delta\delta$ packets, and the rest was left empty. That is, initially there is one big node and the rest are empty, introducing overhead due to token passing through future injections. To evaluate performance, packets were injected at different rates and with different policies. Specifically, we tested injection rates 1 (feasibility upper bound), $1/\sqrt{\delta}$ and $1/\delta$ (approximate theoretical guarantee on delivery rate). Target source nodes for injections were chosen with four different policies: (1) uniformly among source nodes, (2) next (according to ID) source node after the node currently delivering, (3) always in the current source node, and (4) uniformly among all source nodes except the node currently delivering. The idea for policy (3) was to evaluate the system when the token does not circulate, whereas policy (4) can be seen as a worst-case injection since then the likelihood of having one node being big forever is low.

The results of the simulations for the second type of input graph (two overlapped trees obtained running BFS on a connected random graph from different nodes) are illustrated in Figures 1 to 4. Similar results were obtained for the other two types of input graph.

Figures 1, 2, and 3 show the competitive throughput as a function of the number of packets injected for $n = 16, 32$, and 64 respectively. In each of these figures, the 12 combinations of injection rates and policies described above are shown. It can be seen in these plots that the competitive throughput converges to a value $1/(1+\delta)$ when the injection rate is 1, and much higher for smaller injection rates.

In some cases when the injection rate is low, after an initial phase when it converges to high values, the competitive throughput is reduced due to overhead from token passing. Nevertheless, it can be seen in Figures 1 and 2 that still it converges to a value that is $1/(1+\delta)$ or higher. For further illustration of these observations, we include a plot of competitive throughput versus $1+\delta$ in Figure 4. In this plot, the competitive throughput obtained for each combination of network size, injection rate and policy is compared with the $1/(1+\delta)$ lower bound proved in Theorem 3. As in previous plots, it can be seen that for the practical scenarios evaluated, our Multiple-Message Broadcast protocol behaves as shown in our analysis or better.

It is important to notice that the competitive throughput was computed against a theoretical optimal protocol that delivers a packet immediately after injection, which is not possible in practice in a multi-hop network.

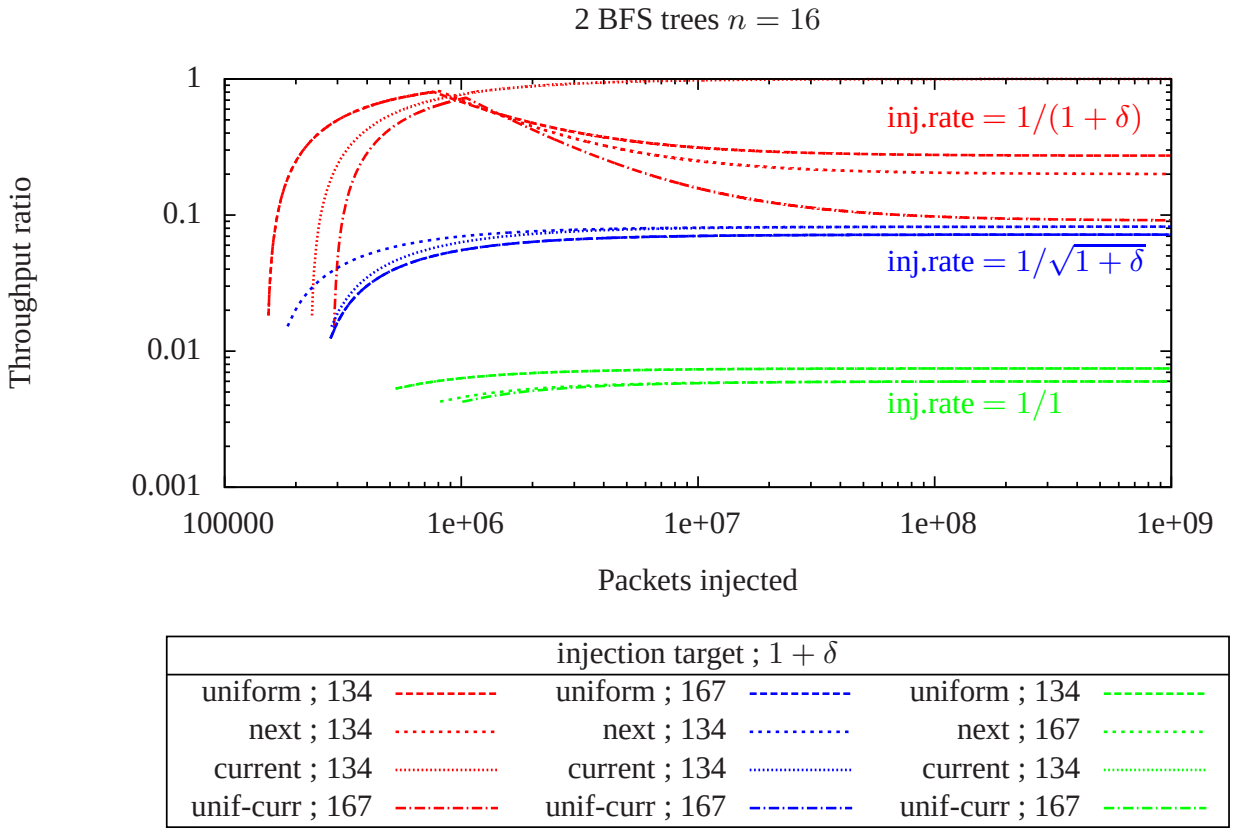

Fig. 1. Competitive throughput vs. packets injected. $n = 16$

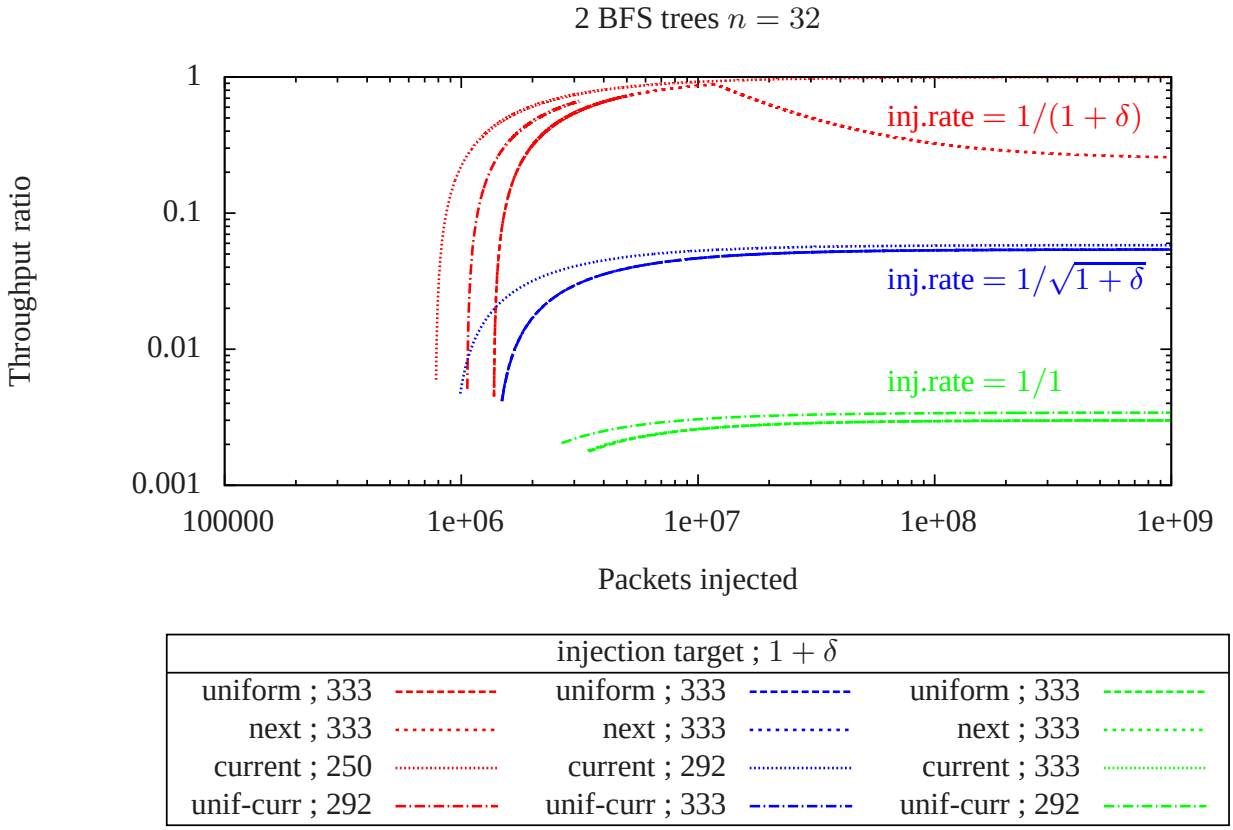

Fig. 2. Competitive throughput vs. packets injected. $n = 32$

References

1. Y. Afek, editor. *Distributed Computing - 27th International Symposium, DISC 2013, Jerusalem, Israel, October 14-18, 2013. Proceedings*, volume 8205 of *Lecture Notes in Computer Science*. Springer, 2013.
2. N. Alon, A. Bar-Noy, N. Linial, and D. Peleg. A lower bound for radio broadcast. *J. Comput. Syst. Sci.*, 43(2):290–298, 1991.
3. L. Anantharamu, B. S. Chlebus, D. R. Kowalski, and M. A. Rokicki. Deterministic broadcast on multiple access channels. In *Proceedings of the 29th IEEE International Conference on Computer Communications (INFOCOM)*, pages 1–5, 2010.
4. R. Bar-Yehuda, A. Israeli, and A. Itai. Multiple communication in multihop radio networks. *SIAM J. Comput.*, 22(4):875–887, 1993.
5. M. A. Bender, M. Farach-Colton, S. He, B. C. Kuszmaul, and C. E. Leiserson. Adversarial contention resolution for simple channels. In *Proceedings of the 17th Annual ACM Symposium on Parallel Algorithms (SPAA)*, pages 325–332, 2005.
6. M. Bienkowski, T. Jurdzinski, M. Korzeniowski, and D. R. Kowalski. Distributed online and stochastic queuing on a multiple access channel. In M. K. Aguilera, editor, *Distributed Computing*, volume 7611 of *Lecture Notes in Computer Science*, pages 121–135. Springer Berlin Heidelberg, 2012.
7. I. Chlamtac and S. Kutten. Tree-based broadcasting in multihop radio networks. *IEEE Trans. Computers*, 36(10):1209–1223, 1987.
8. I. Chlamtac and O. Weinstein. The wave expansion approach to broadcasting in multihop networks. In *Proc. of INFOCOM*, 1987.
9. B. S. Chlebus. *Randomized communication in radio networks*, volume I, pages 401 – 456. Kluwer Academic Publishers, 2001.

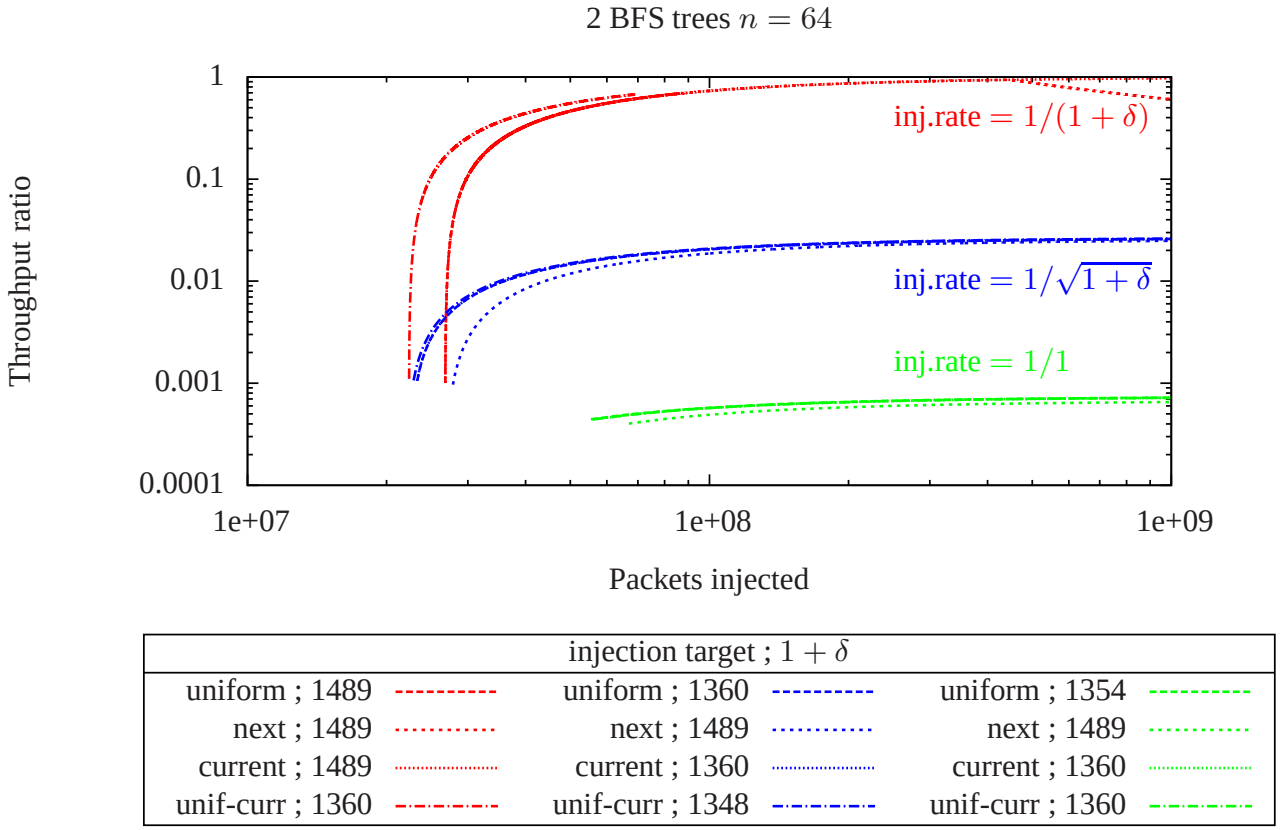

Fig. 3. Competitive throughput vs. packets injected. $n = 64$

10. B. S. Chlebus, D. R. Kowalski, A. Pelc, and M. A. Rokicki. Efficient distributed communication in ad-hoc radio networks. In L. Aceto, M. Henzinger, and J. Sgall, editors, *ICALP (2)*, volume 6756 of *Lecture Notes in Computer Science*, pages 613–624. Springer, 2011.
11. B. S. Chlebus, D. R. Kowalski, and M. A. Rokicki. Maximum throughput of multiple access channels in adversarial environments. *Distributed Computing*, 22(2):93–116, 2009.
12. B. S. Chlebus, D. R. Kowalski, and M. A. Rokicki. Adversarial queuing on the multiple access channel. *ACM Transactions on Algorithms*, 8(1):5, 2012.
13. A. E. F. Clementi, A. Monti, and R. Silvestri. Selective families, superimposed codes, and broadcasting on unknown radio networks. In S. R. Kosaraju, editor, *SODA*, pages 709–718. ACM/SIAM, 2001.
14. A. Czumaj and W. Rytter. Broadcasting algorithms in radio networks with unknown topology. In *FOCS*, pages 492–501. IEEE Computer Society, 2003.
15. S. Daum, S. Gilbert, F. Kuhn, and C. C. Newport. Broadcast in the ad hoc sinr model. In Afek [1], pages 358–372.
16. G. De Marco. Distributed broadcast in unknown radio networks. In S.-H. Teng, editor, *SODA*, pages 208–217. SIAM, 2008.
17. L. Gasieniec, D. Peleg, and Q. Xin. Faster communication in known topology radio networks. *Distributed Computing*, 19(4):289–300, 2007.
18. M. Ghaffari, B. Haeupler, and M. Khabbazian. A bound on the throughput of radio networks. *CoRR*, abs/1302.0264, 2013.
19. M. Ghaffari, B. Haeupler, and M. Khabbazian. Randomized broadcast in radio networks with collision detection. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing*, PODC '13, pages 325–334, New York, NY, USA, 2013. ACM.
20. M. M. Halldórsson and R. Wattenhofer. Wireless communication is in apx. In *Proc. of the 36th International Colloquium on Automata, Languages and Programming, Part I*, pages 525–536, 2009.

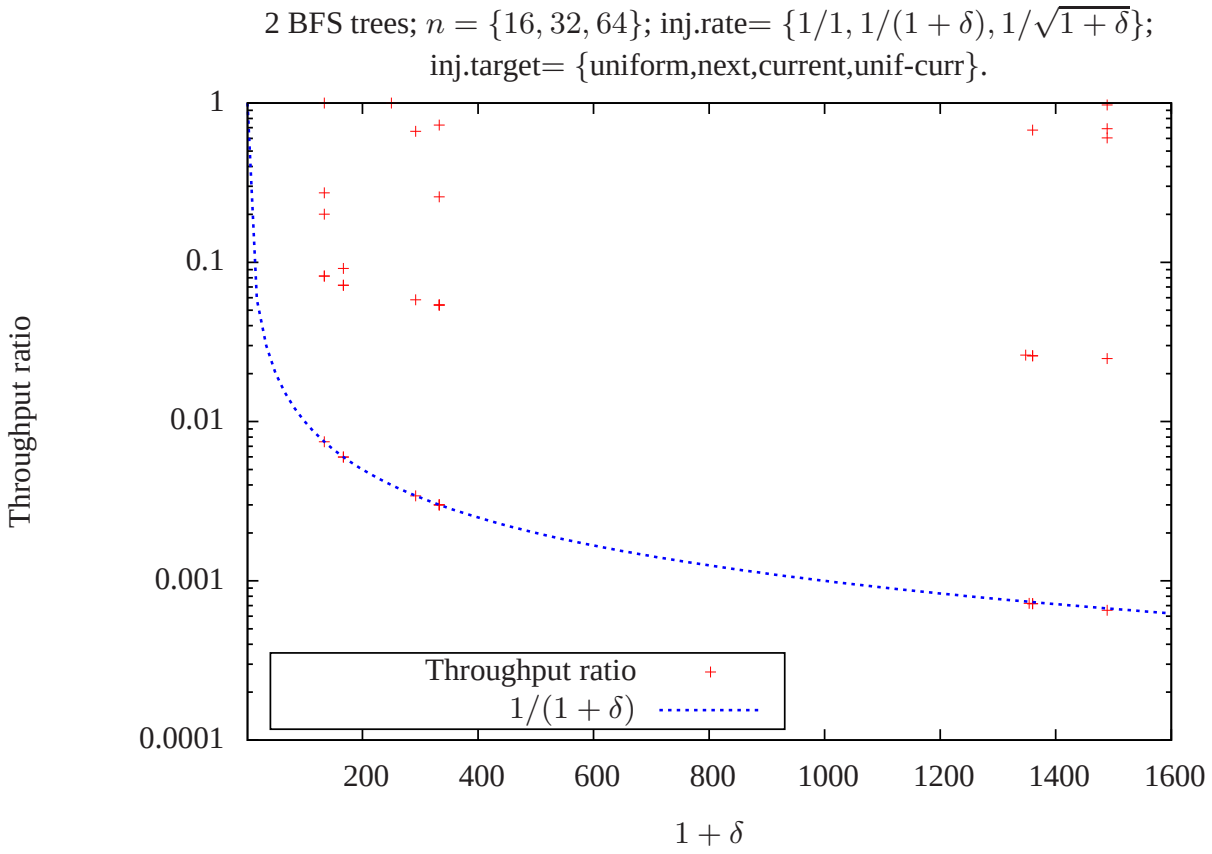

Fig. 4. Competitive throughput vs. $1 + \delta$.

21. T. Jurdzinski and D. R. Kowalski. Distributed backbone structure for algorithms in the sinr model of wireless networks. In M. K. Aguilera, editor, *DISC*, volume 7611 of *Lecture Notes in Computer Science*, pages 106–120. Springer, 2012.
22. T. Jurdzinski, D. R. Kowalski, M. Rozanski, and G. Stachowiak. Distributed randomized broadcasting in wireless networks under the sinr model. In Afek [1], pages 373–387.
23. T. Jurdzinski, D. R. Kowalski, and G. Stachowiak. Distributed deterministic broadcasting in uniform-power ad hoc wireless networks. In L. Gasieniec and F. Wolter, editors, *FCT*, volume 8070 of *Lecture Notes in Computer Science*, pages 195–209. Springer, 2013.
24. T. Kesselheim. Dynamic packet scheduling in wireless networks. In *Proc. of the 31st Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, pages 281–290, 2012.
25. T. Kesselheim and B. Vöcking. Distributed contention resolution in wireless networks. In *Proc. of the 24th International Symposium on Distributed Computing*, volume 6343 of *Lecture Notes in Computer Science*, pages 163–178. Springer-Verlag, Berlin, 2010.
26. M. Khabbazian and D. R. Kowalski. Time-efficient randomized multiple-message broadcast in radio networks. In C. Gavoille and P. Fraigniaud, editors, *PODC*, pages 373–380. ACM, 2011.
27. D. R. Kowalski. On selection problem in radio networks. In *Proceedings of the 24th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 158–166, 2005.
28. D. R. Kowalski, M. A. Mosteiro, and T. Rouse. Dynamic multiple-message broadcast: bounding throughput in the affectance model. In *10th ACM International Workshop on Foundations of Mobile Computing, FOMC 2014, Philadelphia, PA, USA, August 11, 2014*, pages 39–46, 2014.
29. D. R. Kowalski and A. Pelc. Broadcasting in undirected ad hoc radio networks. *Distributed Computing*, 18(1):43–57, 2005.
30. D. R. Kowalski and A. Pelc. Optimal deterministic broadcasting in known topology radio networks. *Distributed Computing*, 19(3):185–195, 2007.

31. E. Kushilevitz and Y. Mansour. An omega(*d* log (*n/d*)) lower bound for broadcast in radio networks. *SIAM J. Comput.*, 27(3):702–712, 1998.
32. C. Scheideler, A. W. Richa, and P. Santi. An o(log n) dominating set protocol for wireless ad-hoc networks under the physical interference model. In *Proceedings of the 9th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 91–100. ACM, 2008.

Appendix

A Modeling SINR under the Affectance Model

Claim. The affectance matrix

$$A(w,(u,v)) = \begin{cases} 0 & \text{if } w = u, \\ \frac{P/d_{wv}^{\alpha}}{P/(\beta' d_{uv}^{\alpha}) - N} & \text{otherwise.} \end{cases}$$

corresponds to the SINR model.

Proof. To prove this claim, we show that there is a successful transmission in the SINR model if and only if there is a successful transmission in the affectance model with matrix A.

Consider a successful transmission in the SINR model. We have

$$\begin{aligned} \frac{P/d_{uv}^{\alpha}}{N + \sum_{w \neq u} P/d_{wv}^{\alpha}} &> \beta' \\ P/(\beta' d_{uv}^{\alpha}) &> N + \sum_{w \neq u} P/d_{wv}^{\alpha} \\ P/(\beta' d_{uv}^{\alpha}) - N &> \sum_{w \neq u} P/d_{wv}^{\alpha}. \end{aligned}$$

If $\sum_{w \neq u} P/d_{wv}^{\alpha} = 0$ then $\sum_{w \neq u} A(w,(u,v)) = 0 \Rightarrow$ success in affectance model. Otherwise, it is $\sum_{w \neq u} P/d_{wv}^{\alpha} > 0$ and we have

$$\frac{P/(\beta' d_{uv}^{\alpha}) - N}{\sum_{w \neq u} P/d_{wv}^{\alpha}} > 1.$$

Thus, it is $P/(\beta' d_{uv}^{\alpha}) - N > 0$ and, hence, we have

$$\frac{\sum_{w \neq u} P/d_{wv}^{\alpha}}{P/(\beta' d_{uv}^{\alpha}) - N} < 1.$$

Therefore, it is $\sum_{w \neq u} A(w,(u,v)) < 1 \Rightarrow$ success in affectance model.

Consider now a non-successful transmission in the SINR model. We have

$$\begin{aligned} \frac{P/d_{uv}^{\alpha}}{N + \sum_{w \neq u} P/d_{wv}^{\alpha}} &\leq \beta' \\ P/(\beta' d_{uv}^{\alpha}) &\leq N + \sum_{w \neq u} P/d_{wv}^{\alpha} \\ P/(\beta' d_{uv}^{\alpha}) - N &\leq \sum_{w \neq u} P/d_{wv}^{\alpha}. \end{aligned}$$

If $P/(\beta' d_{uv}^{\alpha}) \leq N$, it would mean that P is not large enough to overcome the noise, even if no other node transmits. Then, rather than being produced by interference, the failure would be due to consider a link that is not even feasible. That is, $(u, v) \notin E$. Thus, it must be $P/(\beta' d_{uv}^{\alpha}) > N$ and we have

$$\frac{\sum_{w \neq u} P/d_{wv}^{\alpha}}{P/(\beta' d_{uv}^{\alpha}) - N} \geq 1$$

Therefore, it is $\sum_{w \neq u} A(w, (u, v)) \geq 1 \Rightarrow$ failure in affectance model.

B Notes

In this section, we highlight the differences between this paper and our preliminary work appeared in [28].

In [28], we studied a model of affectance that subsumes only *some* SINR models, by combining the effect of Radio Network collisions with affectance from nodes at more than one hop. Here, we generalize our model to subsume any arbitrary interference model. For instance, in the present model it is possible to receive a transmission even when more than one neighboring node transmits, as in some SINR models.

Also, in [28] our maximum path affectance metric was based on fast links only, which yields possibly tighter bounds. However, the definition was based on a specific BFS tree (a GBST [17]) which related the network characterization to our specific algorithmic solution. In the present work the characterization is related only to topology, since it is based on arbitrary BFS trees.

We also notice here that the proof of the maximum rank in [28] has an error, introduced while bounding the maximum number of ranks needed for updating the rank according to affectance. Lemma 1 here provides the correct bound.